A 1-Gigabit Memory System on a Multi-Chip Module for Space Applications

Marianne E. Louie [1], Douglas A. Topliffe [2], Leon Alkalai [1]
[1] Jet Propulsion Laboratory, Pasadena, CA 91109
[2] Sanders, A Lockheed Martin Company, Nashua, NH 03061

Abstract

Current spaceborne applications desire compact, low weight, and high capacity data storage systems along with the additional requirement of radiation tolerance. This paper discusses a memory system on a multi-chip module (MCM) that is designed for space applications. Over 1.5 Gigabits (Gbits) of storage are incorporated into the system to provide 1 Gbit of effective data storage, 0.25 Gbit for either Hamming code error correction check bits or additional data bits if error correction is not desired, and 320 Mbits for fault tolerance and sparing. The memory MCM is based on the Irvine Sensors Corporation 12-high (160 Mbit) DRAM die stacks along with an Actel 1280A Field Programmable Gate Array (FPGA) for the DRAM controller. The design is immune to single event latchup, radiation tolerant to 35 Krads (Si) total dose, and predicted to exhibit a bit error rate << 0.007 errors/day (low earth orbit). A single memory MCM will support a sustained bandwidth of 147 Mbits/sec. Higher bandwidth can be achieved in a system by incorporating additional memory MCMs.

1 Introduction

Spacecraft trends are moving toward smaller, lightweight designs to reduce overall life-cycle costs. However, mission trends toward spacecraft autonomy and on-board science data analysis are necessitating higher memory capacities, thus creating a conflicting requirement to smaller mass and volume in space-based solid-state recorder (SSR) subsystems. By eliminating chip packaging bulk through multi-chip module (MCM) technology and die stacking, both an increase in memory capacity as well as a reduction in mass and volume can be achieved. The 1-Gigabit mass memory MCM assembly, jointly developed between Sanders and the NASA Jet Propulsion Laboratory, directly addresses the need for compact, space qualifiable memory systems. The use of 3D memory die stacks and a Field Programmable Gate Array (FPGA) based controller, integrated onto a single MCM, provides a high density building block for data storage system needs. The controller design provides Hamming code error detection and correction, DRAM sparing, and block sequential memory accessing to support the needs of fault tolerant spaceborne memory systems. Some potential applications for the memory MCM include (1) the solid state recorder for the Pluto Express mission and (2) the Sanders Solid State Recorder.

1.1 Spacecraft Storage Technology and Trends

In the last several decades, the National Aeronautics and Space Administration (NASA) has launched a fleet of highly capable robotic spacecraft for the exploration of deep space. Missions such as Voyager 1, Voyager 2, Magellan, Ulysses, and Galileo have been extremely successful in furthering mankind's understanding of our Solar System. These missions have generally been very complex, and thus required many years to design, develop, and operate. Moreover, the cost for each of these programs generally exceeded one billion dollars during the mission's life-cycle. Programs that spanned ten years for design and development would typically have obsolete technology by the time of launch.

For the past several years, NASA's declining budget has initiated a set of ambitious, highly focused, and yet more frequent and less costly missions to explore both deep space and our own planet Earth. Consider as an example the memory technology on the Galileo spacecraft which was launched in October 1989. The Command and Data Handling subsystem has a total of 384 KBytes of static random access memory (SRAM) using a total of 864 integrated circuit chips (ICs) where each IC is 4 Kbits; and the Attitude and Articulation Control Subsystem (AACS) has a total

of 128 KBytes of memory (SRAM) using more than 1000 ICs where each IC is 1 Kbit. Contrast this to the Mars Pathfinder (low-cost) mission scheduled to be launched in December 1996 in which a total of 64 16-Mbit DRAM chips implements a 1 Gbit storage device for the single board spacecraft computer.

As shown in the table below from the recently published National Technology Roadmap for Semiconductors [4], new generations of DRAM devices are expected every three years, driven by advances in the optical lithography technology that enables further reduction in the minimum feature size.

Year	Min. Feature Size	DRAM Capacity
1995	0.5	16 Mbit
1998	0.35	64 Mbit
2001	0.25	256 Mbit
2004	0.18	1 Gbit
2007	0.12	4 Gbit
2010	0.08	16 Gbit

In the design described in this paper currently available yet advanced 16 Mbit DRAM die stacks are used to implement a 1 Gbit solid state recorder for spaceborne applications. Additional use of Multichip Module technology enables integration of over 100 IC's into a single module which results in a highly compact, light-weight, and reliable solid state recorder product that weighs on the order of 100 grams. This design is also adaptable to future memory and stacking advances. Missions such as the NASA Pluto Express project which starts in 1999 and launches in 2001 may, in fact, plan to use 64-Mbit DRAMs to implement a 4 Gbit solid state recorder within a mass budget of 100 grams. This in itself is a major enabling technology for many future missions of the New Millennium.

2 Design Overview

The mass memory module presented in this paper is designed as a 32 Mbit x 40 easy-to-use building block for space-based solid-state recorder (SSR) systems. The module uses dynamic random access memory (DRAM) which provides a high memory density to mass ratio. The module's internal memory controller handles the required DRAM refreshing, and DRAM control so that the memory module can be easily used much like static random access memory (SRAM). Optional Hamming code error detection and correction (EDAC) is provided such that the 40-bit data width can support either (1) a 32-bit data word with 7 bits of single bit error correcting or double bit error detecting Hamming code check bits or (2) a 40-bit data word without Hamming code EDAC. Error detection and correction is handled internally by the controller. Error diagnostics, however, are externally available to allow external activation of the DRAM spares according to each SSR system's individual requirements. When the internally supported EDAC is used the effective data capacity of the memory MCM is 1 Gbit. The effective data capacity increases to 1.25 Gbits when the module is used as a 40-bit wide data storage without EDAC.

Figure 1 provides a high level description of the module and its address, data and control lines. Accessing the module's memory is simply controlled by the MCM enable, output enable (OE), read/write control (R/W), address and data lines. A 40-bit data width accommodates 32 bits of data and either (1) an additional 8 bits of data or (2) up to 8 check bits for EDAC. The remaining control and output lines include:

- Blk_Enable — initiate reading or writing of sequences of blocks with sequential addressing. An initial block address and block count must be specified through the address and data lines.
- Refresh_Scrub — initiate scrubbing during the refresh process (detection and correction of errors within memory) versus refresh-only without scrubbing.
- EDAC_Bypass — the EDAC logic can be bypassed for applications that do not want EDAC processing or prefer a 40-bit data width.
- Display_Err — display the address of the last EDAC error; the address specifies the 4-word block address, memory bank and die level within the 3D stack of DRAM die.
- Reconfig — reconfigure the internal address map to replace failed memory dice with spares.
- EDAC_Err — output status line indicating the detection of an error during EDAC.
- Corr_Err — output status line indicating the correction of a single bit error in a 32-bit word during EDAC.
- MCM_Busy — output status line to indicate that the MCM is currently busy processing. To conserve registers, commands are not queued within the memory module.

- Accessing_Blks — output status line indicating that the MCM is currently accessing a sequence of blocks. This line can also be monitored to determine when refreshing is taking place before/between a series of sequential block accesses.

Additional lines are provided for testability, including direct access to the DRAM and controller component lines.

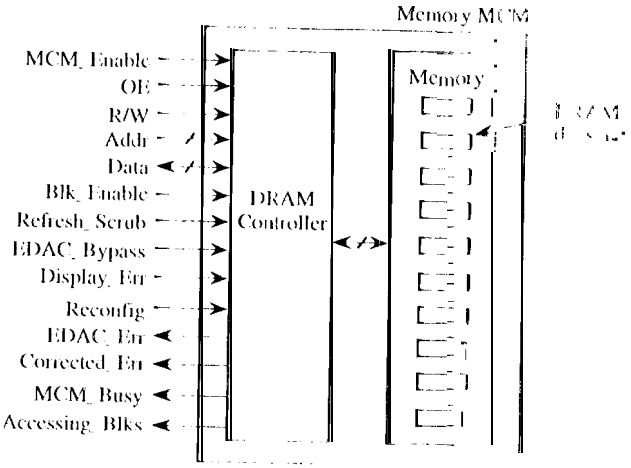

Figure 1: Memory MCM Overview

The DRAM controller design is based on the Actel 1280A Field Programmable Gate Array (FPGA) [1]. The flexibility of an FPGA enables faster and cheaper breadboarding and prototyping than custom Application-Specific Integrated Circuits (ASICs). An FPGA also allows a wide variety of controller speeds, refreshing schemes, and custom controller schemes to be offered in a final product line. FPGAs, however, are I/O pin limited; as a result the memory MCM design uses a two-phase row/column-like approach to asserting the memory address.

The memory module is designed with the Irvine Sensors Corporation (ISC) 12-high DRAM stacks [2] where each stack yields 10 usable die. Each stack provides 160 Mbits of storage in which one 16-Mbit IBM Luna-ES die is accessible at any given time along a 4-bit wide data path. On the ISC 12-high DRAM stacks, the corresponding data lines, address lines and column address strobe (CAS) of each Luna-ES die are tied together while the row address strobe (RAS) lines are independent.

3 Memory Partitioning

Radiation testing of some IBM Luna series die (Luna-E and Luna-C) has shown that errors can occur on entire rows and columns of a DRAM die [3]. To avoid uncorrectable multiple errors by any single DRAM die failure, the memory module places no more than one bit of any given word on a single Luna-ES DRAM die. Each DRAM die, however, has a 4-bit wide data path; the memory module thus implements bit interleaving of 4 words for efficient utilization of the memory space (Figure 2). As a result, each single memory read or write accesses four sequential words - the definition of a single 4-word block on the memory module.

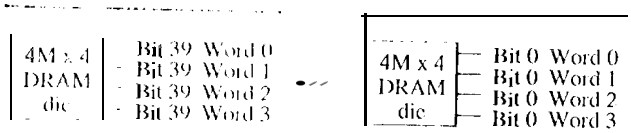

Figure 2: Interleaving of Four Words

Only 23 address bits are required to access the 1.25 Gbits of data and check bits in the memory MCM because each single address references a 160-bit single 4-word block. A two phase row/column-like addressing approach uses 13 address lines to assert the 23-bit block address. In the first phase, thirteen address bits specify both (1) the memory bank to be accessed and (2) the row address to be accessed in each of the active die in the bank. At the second phase, ten address bits specify the column of memory to be accessed in each DRAM die.

Regarding memory capacity, ten ISC 12-high DRAM stacks (Figure 3) provide 1.25 Gbits for data storage and EDAC check bits plus an additional 320 Mbits for sparing. The upper and lower portions of the stacks are divided into two separate memory banks where the corresponding subbanks of each bank (half a bank) share two sets of spares (Figures 3 and 4).

Figure 4 shows the organization of four interleaved words within the two subbanks of a 0.5 Gbit memory bank. Each four-word block access requires four 40-bit accesses to the memory bank where (1) the first access is to logical die level 0 of the subbank and reaches the least significant ten bits of each of the four words, (2) the second access is to logical die level 1 and reaches the second least significant bits of each of the four words, (3) the third access is to logical die level 2 and reaches the next ten more significant bits, and (4) the last access reaches the ten most significant bits of all four words. Hence, four active dice of each DRAM

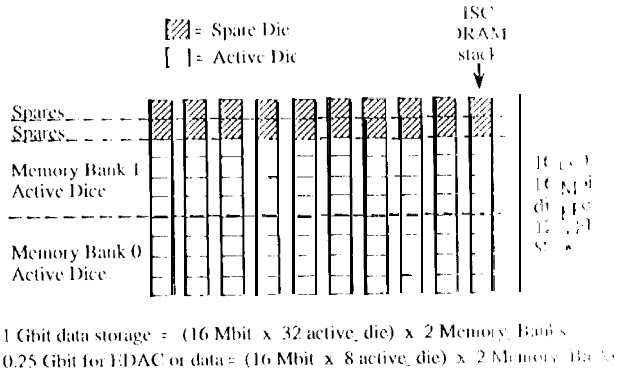

Figure 3: Memory Partitioning in the DRAM Stacks

stack, corresponding to the given memory bank are accessed during a read or write of a data block. The control sequences to all four die levels are pipelined to reduce the total access time of the 4-word block.

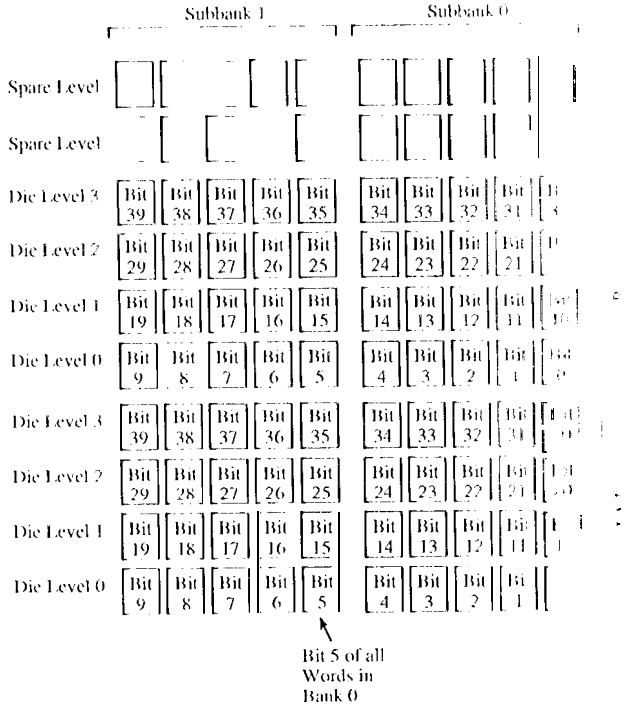

Figure 4: Bit Organization

4 DRAM Access Modes

The ISC DRAM stacks offer a variety of access modes including, [2]: read, early write and delayed write, read-modify-write, fast page mode read and write, fast page mode read-modify-write, RAS-only refresh, CAS-before-RAS refresh, and hidden refresh read and write. With the memory MCM design, the use of bit interleaving (Fig. 2) restricts the selection of DRAM access modes for the controller. For example, while DMA-like memory schemes use fast page mode to rapidly pipeline sequential word accesses, an access to four sequential words (block) within the memory MCM requires four accesses to separate die levels rather than to sequential words within a single die level. Despite the limited DRAM accessing operations, a bandwidth of 147 Mbits/sec can be maintained for sequences of blocks; this is achieved through pipelining within the DRAM accessing control sequence.

Selection/elimination of DRAM access modes is also based on a simplification of the DRAM controller. For example, RAS-only refreshing enables flexible control over the refreshing counter so that bit interleaved words can be refreshed and scrubbed with common hardware logic. And, early write is chosen over delayed write due to the simplified logic and logic commonality with a read operation. As a result, the selected DRAM access modes for the memory MCM design is reduced to read, early write, and RAS-only refresh — this selection choice simplifies the controller design as well as satisfying the constraints imposed by bit interleaving.

5 Memory System Packaging Approach

5.1 Requirements

The 1-Gbit mass memory system is intended for long duration spaceflight applications. Memory system development therefore hinges on selecting an appropriate MCM package to satisfy the requirements for space qualification and space missions. Among the mandatory MCM packaging attributes are high reliability, low weight and small size — features that are important to the trend of smaller and better spacecraft. A summary of key attributes for a space-qualified MCM product/prototype include:

- Reliability

- Low weight and small size
- Quantifiable performance capabilities
- Path to Class H and K quality screening levels
- Reproducible product/process
- Low recurring costs

Additional packaging requirements specific to the 1-Gbit memory MCM include:

- Package body dimensions not to exceed 2" x 2"
- Weight not to exceed 100 grams
- Leaded surface mount package
- 300 leads minimum, 25 mil pitch minimum
- Power dissipation > 4W continuous
- Shock, 50 G's @ 1 KHz - 10 KHz
- Hermetically sealed package
- Outgassing: < 1% Total Weight Loss (TWL) < 0.1% Volatile Condensible Material (VCM)
- Accommodations for memory die stacks

To date, numerous MCM technologies have been investigated and evaluated for the mass memory system. These include:

- Packaging
  - Aluminum Nitride/ Ceramic (hermetic package)
  - Printed Wiring Board (PWB)
  - Plastic encapsulation
- Substrate/Interconnection
  - Multi-layer (ceramic/PWB)
  - Silicon substrate
  - High Density Interconnect (HDI), a General Electric developed thin film polyimide overlay technology
- Component Connection
  - Wire bonding
  - HDI

Special packaging constraints imposed by the memory die stacks have also been identified and evaluated

5.2 Packaging Challenges from Memory Die Stacks

Memory die stacks present several unique packaging challenges which are understood and uniquely resolvable by each MCM manufacturer. The issues associated with the stacks include stack height, variation in stack height, and configuration of the stacks. First, stack height significantly influences the minimum separation/spacing between stacks on the MCM. For tall stacks (190+ mils), sufficient spacing must be provided to allow the automatic wire bonding equipment to reach between the stacks for bonding to the substrate. Mechanical requirements on wirebonds, such as aspect ratio, also affect stack spacing. The net result of these two requirements is a forced separation of adjacent die stacks.

HDI interconnect technology does not present the stack spacing issues of wirebonding. However, height variations rather than overall stack height impose additional processing steps. The HDI process assumes that all stacks are almost exactly the same height. When stack height variations exceed the HDI tolerated height variations, a unique baseplate milling/shimming for each stack is required to achieve uniform connection plane elevation.

Variations in usable die configuration within a stack also present a challenge to MCM fabrication. For yield improvement, the memory die stacks have two spare die so that any ten of the twelve stacked die are guaranteed to be known-good when delivered to the customer. As a consequence, each stack's usable die configuration can be different. Through selective wire bonding/connection, the MCM assembly process must account for these configuration differences, thus entailing an extra step in the fabrication process.

5.3 MCM Technology Evaluation

Of the evaluated packaging technologies, the baselined Ceramic/Aluminum Nitride package with wire bonded die connection provides the most mature production path to the MCM with moderate recurring costs. However, the limitations imposed by wire bonding result in varying degrees of memory stack packing efficiencies and can result in a larger overall package with increased weight. The potentially long wire bond connection also creates electrical properties that may require special attention.

Two advanced substrate and electrical connection approaches to MCMs were also examined; these include a PWB with chip (bare die)-on-board, and plastic encapsulation with HDI. Both approaches address

the reduction of MCM size, and hence overall weight. The use of alternative materials also helps to reduce weight. Due to the interconnect technology of HDI and a multi-level interconnection capability within the combination of PWB and HDI, die/stack/component density (spacing) is effectively higher. Additionally, the PWB and plastic encapsulation approach provides a 10I% mass, integral package body for the MCM. lead frames 0 1 " flex-print can be used for input, and output of signals to/ from the MCM. For the second advanced approaches, an additional package with less associated weight is needed for hermeticity.

Each of the evaluated packaging and interconnection approaches has demonstrated desirable features. We are currently in the process of selecting the best approach for our needs.

6 Summary

This paper has presented the high-level design of a 1-Gbit single error correcting memory system, in an MCM (or] .25-Gbit non-error checking memory) for spaceborne applications. Ten Irvine Sensors 16 high 160 Mbit DRAM stacks provide the required storage for both data, EDAC, and spares. While die stacking presents height, height variation and die configuration issues to MCM fabrication, MCM technologies provide a wide variety of packaging options to resolve these issues. Together, die stacking and MCM technology help to achieve the goal of reduced weight and size for space applications.

The presented memory module design is immune to single event latchup, radiation tolerant to 35 krads (Si) total dose, and is predicted to exhibit a bit error 101" rate << ()-()()2? errors/day (low earth orbit). The memory MCM design provides easy-to use features such as internal refresh management, scrubbing versus refresh-only options, block sequential access, no EDAC and an EDAC-bypass option, and simple accessing control similar to that of accessing SRAM. Both small and large spaceborne data storage systems ('0111(1 be made smaller and lighter with the memory module as a building block.

Acknowledgments

The work described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, and by Sanders, a Lockheed Martin Company. This work was sponsored by the National Aeronautics and Space Administration and by the Lockheed Martin Company. We would like to thank Dr. Carl Kukkonen and 1)1 Virendra Sarohia of the Center for Space Microelectronics Technology at the Jet Propulsion Laboratory and Dr Gordon Johnston of the National Aeronautics and Space Administration for their partial sponsorship of this work under the Advanced Flight Computing Task at the Jet Propulsion Laboratory.

References

[1] Actel FPGA Data Book and Design Guide, 1994.

[2] Irvine Sensors Corp. Product Application Specification 40 Meg x 4 Bits DRAM Stack Industrial Screening, Preliminary Data Sheet, April 1996.

[3] P. Calvel, et al., "Space Radiation Evaluation of 16 MBit DRAMs for Mass Memory Applications," Trans. of the Nuclear Radiation Effects Conference, December 1994.

[4] The National Technology Roadmap for Semiconductors, Semiconductor Industry Association, December 1994.